

# Report on Disc Tuning

J. Campbell


University of Birmingham

21st August 1998

Aim To improve the data transfer rate to discs on SUN 450 and to demonstrate a minimum 20 MB/sec write performance


## Initial setup

On first examination, the system was found to have two 9 GB discs (in addition to a system disc and swap area) each configured with three partitions:



S0 is a small 2 MB partition for the SDS (Solstice DiskSuite) database. S1 and S6 are two ~4 GB partitions.

D0 and D1 are meta-disks in the SDS system and were configured as:




With this configuration the second 4 GB partition on D0 was written only when the first was full. Similarly for D1. Thus the slices are accessed serially and not in parallel. In other words, maximum throughput is limited to single disc performance.

The performance measured for Write/Read was 11.42/14.11 MB/sec. The transfer size was 500 MB.

#### Revised Configuration: 1

The SDS configuration was then revised so that meta-disk D0 consisted of two partitions in parallel:




The performance measured for Write/Read was 18.62/25.60 MB/sec. I/O statistics show that this rate was obtained by parallel read rates of 13.04 MB/s on each partition. File transfer size was 500 MB.

Note that both discs are on the same SCSI controller (c2).

#### Revised Configuration: 2

A third disc was added on another SCSI controller (c3) to obtain the final configuration:



The performance measured for in this configuration for Write/Read was 25.74/46.43 MB/sec for a file transfer size of 500 MB.

When this measurement was repeated with 1 GB file transfer size, the Write/Read performance was 25.31/36.0 MB/sec. I/O statistics show that this rate was obtained by parallel read rates of 12.3 MB/s on each partition.

### Conclusion

With proper configuration of the meta-disks, it will be straightforward to obtain write data rates in excess of 20 MB/sec. The different write/read rates are due, in part, to the nature of the discs. This difference can be seen on the disc specifications.

### Recommendation

It is inadvisable for performance reasons to partition the discs. These then form two separate discs and cannot be written in parallel. The small SDS database partition is excluded from this statement.

For maximum throughput, no more than 3 discs should be on one controller. Thus if three Ultra-SCSI controllers are available, 3 (SCSI) x 3 (discs) x 8 GB/disc provides a total of approximately 72 GB of disc space with a potential performance of write/read for 1 GB files of 25/36 MB/sec. This is within the 40 MHz limit of Ultra-SCSI.